(54) Title: VIDEO DATA TRANSMISSION METHOD AND APPARATUS
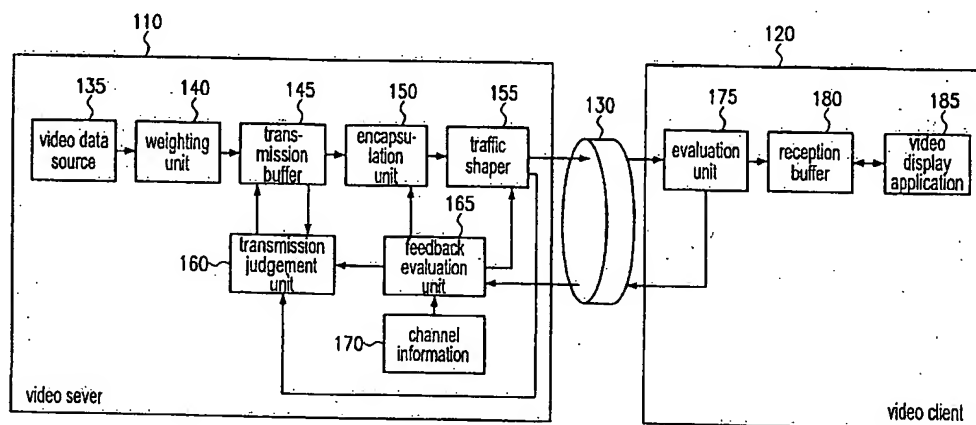
(57) Abstract: An apparatus and a method are provided for transmitting video data to a video client. The video data is a series of video data units. A unique priority level is assigned to each video data unit that is to be transmitted. The video data units are then buffered, wherein a first part of the video data units have not yet been transmitted, and a second part of the video data units are available for retransmission. According to the embodiments, one of the buffered video data units is selected from the first and second parts, according to its priority level, and is transmitted to the video client. The embodiment allows for improving the video quality in particular when transmitting compressed video streams over unreliable links, by scheduling regular transmissions as well as retransmissions using priorities.

VIDEO DATA TRANSMISSION METHOD AND APPARATUS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to an apparatus and method for transmitting video data to a video client and is in particular applicable for compressed video streaming of e.g. MPEG-4 data over unreliable links such as wireless channels.

2. Description of the Related Art

Generally, a video stream consists of a flow of bit packets called frames that contain information about the video picture itself as well as some meta information. The meta information includes for instance the time-to-send, i.e. the time at which the frame shall be sent, and the time-to-display, i.e. the time at which the frame shall be displayed.

The frames contain the video information in different formats, depending on the video compression. Some of these frames can be used independently while others dependent on the preceding frames. The independent frames are called intra-coded frames or I-frames, whereas the dependent frames are so-called predictive-coded frames or P-frames. These dependencies lead to a different importance of the frames. The importance of the P-frames decreases with the separation to the preceding I-frame.

Two main problems arise when transmitting real time video streams over unreliable links.

The first problem is that transmitting video frames over an unreliable channel leads to packet loss. Compressed video streams are however extremely susceptible to packet loss because of the inter-dependencies of the video frames.

For that reason, several mechanisms have been developed for reducing the packet loss and for compensating the consequences thereof. One common technique is to retransmit lost packets by simply signaling the loss of a packet from the receiver to the transmitter and letting the transmitter repeat

5   the lost packets. This is called Automatic Repeat Request (ARQ).

The second problem that arises when transmitting real time video streams over unreliable links is a consequence of the delay requirements and/or limited bandwidth of the transport channel. All video frames have a time-to-send and a time-to-display value. Frames that are received after their time-

10   to-display time, are of no use at the receiver. Thus, frames are sometimes delayed by the channel to an extent that they are useless.

To control the display times and cope with real time requirements, the Real-time Transport Protocol (RTP) was developed. This protocol adds some information such as time stamps and sequence numbers to the video frames

15   that are encapsulated in RTP packets. This information is used to send the video frames at the scheduled time to the right video application. Some measurements such as delay jitter, simple loss statistics, etc can be carried out at the receiver, and the results are signaled to the source by means of the RTP control protocol (RTCP).

20   Considering the above-mentioned retransmission techniques, because of the delay requirements not all of the lost packets can be retransmitted. Further, as different data packets are of different importance, just simply retransmitting lost packets without any criterion would cause less important packets to block the transmission of other more important packets, leading to

25   a decrease in the quality of the received video stream.

EP 1 130 839 A1 discloses a data transmission method and apparatus that improve the received quality by retransmitting the I-frames as well as a number of subsequent P-frames, with the number set according to the current channel conditions. That is, the disclosed technique does not limit

30   the retransmission to the I-frames or to the I-frames and a fixed number of P-

frames, but is adaptive to the channel conditions. This allows for more efficiently using the available channel bandwidth, but there are still limitations that would allow for even better improving the quality of the video streaming application while meeting the time requirements and channel bandwidth
5   limitations.

SUMMARY OF THE INVENTION

An improved apparatus and method for transmitting video data to a video client are provides that may efficiently use the available bandwidth and lead to improved video quality.

10   In one embodiment, there is provided an apparatus for transmitting video data to a video client. The video data is a series of video data units. The apparatus comprises a priority assignment unit for assigning a unique priority level to each video data unit to be transmitted. The apparatus further comprises a transmission buffer for buffering a plurality of video data units
15   each having assigned a priority level. A first part of the plurality of video data units includes video data units that have not yet been transmitted. A second part of the plurality of video data units includes video data units that are available for retransmission. The apparatus further comprises a transmission judgement unit for selecting, from said first and second parts, one of the
20   buffered video data units according to its priority level. The apparatus further comprises transmission means for transmitting the selected video data unit to the video client.

In another embodiment, a method of transmitting video data to a video client is provided. The video data is a series of video data units. The method
25   comprises assigning a unique priority level to each video data unit to be transmitted, buffering one or more video data units each having assigned a priority level, wherein a first part of the video data units have not yet been transmitted and a second part of the video data units are available for retransmission, selecting one of the buffered video data units from the first

and second parts according to its priority level, and transmitting the selected video data unit to the video client.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings are incorporated into and form a part of the

5 specification for the purpose of explaining the principles of the invention. The drawings are not to be construed as limiting the invention to only the illustrated and described examples of how the invention can be made and used. Further features and advantages will become apparent from the following and more particular description of the invention, as illustrated in the

10 accompanying drawings, wherein:

FIG. 1 is a block diagram illustrating the components of a video server and video client that are operating according to an embodiment;

FIG. 2 is a graph illustrating the priority level assignment within a sequence according to an embodiment; and

15 FIG. 3 is a graph illustrating the priority level assignment to data units of multiple sequences according to another embodiment.

DETAILED DESCRIPTION OF THE INVENTION

The illustrative embodiments of the present invention will be described with reference to the figure drawings wherein like elements and structures are

20 indicated by like reference numbers.

Referring now to FIG. 1, a video server 110 is depicted that is connected via an unreliable channel 130 to a video client 120. The channel 130 is considered unreliable whenever it produces a propagation delay and/or bit errors. Examples of unreliable channels are wireless links.

25 The video server 110 transmits video data through the channel 130 to the video client 120. The video information is provided in the form of frames or frame segments by the video data source 135 of the video server 110. The

frames or the segments of a frame are encapsulated in Packet Data Units (PDU). A PDU may contain a single frame or a segment of a frame.

The PDUs are forwarded to the weighting unit 140 where a priority value, or weight, is assigned using some PDU meta information such as the expiration
5   time, the type of compression and the sequence number of the PDU. The expiration time is the time at which a PDU can be sent at the latest, i.e. the time difference of time-to-display and channel propagation delay.

The PDUs are then stored in the transmission buffer 145. The transmission buffer 145 is controlled by a control unit (not shown) that manages the
10   expiration times of each PDU in the transmission buffer 145 and deletes those PDUs that have been expired. The transmission buffer 145 is further controlled by the transmission judgment unit 160 that decides which PDU is to be sent, resent or discarded by selecting the respective PDU according to its priority level (or weight). Once the transmission judgment unit 160 has
15   selected a PDU for (re)transmission, the PDU is forwarded to the encapsulation unit 150 that encapsulates the respective PDU into an RTP PDU by adding appropriate header information. The header information of the RTP PDU is different for transmissions and retransmissions.

The RTP PDU is then forwarded to the traffic shaper 155 which sends the
20   RTP PDU at the bit rate given by the channel information block 170.

There are several possibilities for triggering the transmission judgment unit 160 to select the highest priority PDU in the transmission buffer 145.

One possibility is that the transmission judgment unit 160 is triggered by the feedback evaluation unit 165 that receives and evaluates any feedback that
25   is received from the video client 120. The feedback from the video client 120 may include a request for retransmissions or an acknowledgment indicating that a PDU was received correctly. The video server 110 then has to decide whether to send a retransmission or whether to delete frames from the transmission buffer 145 that have been acknowledged.

Another possibility of triggering the transmission judgment unit 160 is to use a control signal from the traffic shaper 155. The traffic shaper 155 signals the completion of each PDU transmission to the transmission judgment unit 160, indicating that the video server 110 is now ready to continue sending
5   PDUs from the transmission buffer 145.

A further possibility is that the transmission judgment unit 160 is triggered by the transmission buffer 145 itself. When the transmission buffer 145 is empty, any incoming PDU can be immediately transmitted.

While three different possibilities of how to trigger the transmission judgment
10 unit 160 have been described, it will be apparent to those of ordinary skill in the art that other possibilities might exist. Further, it is apparent that more than one triggering scheme can be employed simultaneously.

At the receiving side, PDUs are received that have undergone propagation delay and bit errors by passing the unreliable channel 130. In the evaluation
15 unit 175, the header information of each received RTP PDU is evaluated, the receiving statistics are updated and, if needed, retransmission requests and acknowledgments are generated. The generated messages include ACK (acknowledgment), NACK (non-acknowledgment) or MACK (multi-purpose acknowledgment) message and are transmitted back to the video server 110.

20 The evaluated PDUs are then forwarded to the reception buffer 180 where the header is stripped off and the frame is stored. The video display application 185 will then manage to release the frames from the reception buffer 180 according to the display times and carry on the actual display of the video information.

25 According to the embodiment, each PDU is assigned a unique priority level. For this purpose, each frame of a sequence is classified into one of several priority groups, or classes. A sequence is a series of frames starting with an I-frame and further incorporating all P-frames that depend on this I-frame. Frames or segments of a frame are encapsulated in PDUs. Each PDU is
30 then assigned a unique weight (priority) that depends on its priority group.

A priority group is a group of PDUs that belong to certain types of frames. In the present embodiment, three priority groups (or types of frames) are distinguished: I-frames (I), more important P-frames (P_m) and less important P-frames (P_i). The weights of the PDUs that belong to the different priority groups are calculated according to the following equations:

$$\hat{w}_I = w_{max} - G \cdot D \cdot SQN$$

$$\hat{w}_{Pm} = \hat{w}_I - D \cdot (1 + G \cdot L)$$

$$\hat{w}_{Pl} = \hat{w}_{Pm} - D \cdot (1 + G) = \hat{w}_I - 2 \cdot D - D \cdot G \cdot (1 + L)$$

In these equations, G denotes the number of priority groups, i.e. $G = 3$ in the present embodiment.

D is the maximum number of PDUs in the groups:

$$D = \max\left\{ I_{max}, P_{m\_max}, P_{l\_max} \right\}$$

where I_{max} is the maximum number of PDUs in the I-frame priority group, P_{m_max} the maximum number of PDUs in the more important P-frame priority group, and P_{l_max} is the maximum number of PDUs in the less important P-frame priority group. These numbers are measured per sequence for all sequences.

SQN is the number of a sequence. As stated above a sequence is a series of PDUs starting with an I-frame and further including all P-frames that depend on this I-frame.

Finally, the parameter L is a value indicating the length of the weighting interval. The weighting interval will be explained in more detail below.

Referring now to FIG. 2, the graph illustrates, in an example, the dependencies of the above shown equations. The graph of FIG. 2 is given for explanatory reasons only, to illustrate the concept of providing priority groups.

As the non-fragmented frames happen to be very large, they are fragmented to achieve smaller PDU sizes for efficient transmission. Additionally, the fragmentation is used to obtain substantially constant PDU sizes, as far as this is possible. Considering all sequences, I-frames are segmented to a maximum of three PDUs, P_m-frames are segmented to a maximum two PDUs and P_l-frames are also divided to a maximum of two PDUs.

In the example of FIG. 2, a video stream includes in each sequence one I-frame and five successive P-frames. The first three of the five P-frames belong to the more important P-frame priority P_m, and the remaining two P-frames belong to the less important P-frame priority group P_l.

This results to the following maximum number of PDUs in each priority group:

$$I_{max} \quad = 3 \cdot 1 \ = 3$$

$$P_{m\_max} \ = 2 \cdot 3 \ = 6$$

$$P_{l\_max} \ = 2 \cdot 2 \ = 4$$

Thus, by applying the equation that defines the D parameter, the maximum number of PDUs for all priority groups is $D = 6$ in the present example.

In FIG. 2 the PDUs in each priority group are shown to have the same weight i.e. priority level. It will however be appreciated by those of ordinary skill in the art that it is rather preferred that each PDU has its unique weight. An embodiment will therefore now be discussed in more detail with reference to FIG. 3.

In this example, the priority values of each PDUs in a number of subsequent sequences are shown to gradually decrease. In the example of FIG. 3, the maximum weight w_{max} is set to a constant value of 65,535 which corresponds to a 16bit address. The individual priority values are then assigned according to the following equations:

$$w_l = \hat{w}_l - PN$$

$$w_{P_m} = \hat{w}_{P_m} - PN$$

$$w_{P_l} = \hat{w}_{P_l} - PN$$

In these equations, the parameter PN denotes the number of a PDU in a priority group. At the beginning of each priority group, this number is zero. By applying these equations, the priority levels that are assigned to the PDUs in each sequence are:

PDU	SEQ=0	SEQ=1	SEQ=2	SEQ=3	SEQ=4	SEQ=5	SEQ=6	SEQ=7	SEQ=8
0	65535	65517	65499	65481	65463	65445	65427	65409	65391
1	65534	65516	65498	65480	65462	65444	65426	65408	65390
2	65533	65515	65497	65479	65461	65443	65425	65407	65389
3	65475	65457	65439	65421	65403	65385	65367	65349	65385
4	65474	65456	65438	65420	65402	65384	65366	65348	65384
5	65473	65455	65437	65419	65401	65383	65365	65347	65383
6	65472	65454	65436	65418	65400	65382	65364	65346	65382
7	65471	65453	65435	65417	65399	65381	65363	65345	65381
8	65470	65452	65434	65416	65398	65380	65362	65344	65380
9	65451	65433	65415	65397	65379	65361	65343	65325	65379
10	65450	65432	65414	65396	65378	65360	65342	65324	65378
11	65449	65431	65413	65395	65377	65359	65341	65323	65377
12	65448	65430	65412	65394	65376	65358	65340	65322	65376

The priority level assignment according to the embodiment is intended for giving lower weights to less important PDUs to the benefit of the more important ones. Assuming that for each sequence the PDUs of this sequence expire after five subsequent sequences, the expiration time is set to a value that corresponds to the length of five sequences. For instance, the PDUs of the first sequence ($SQN = 0$) expire before the sixth sequence ($SQN = 5$) starts. In the present embodiment, the PDUs of the P_m frame in the first sequence are made less important than all I-frame PDUs of the first to fourth sequences. This is shown in FIG. 3 by means of the horizontal dashed line.

This interval of sequences will be called weighting interval in the following description. The weighting interval is therefore the interval in which the PDUs of the more important P_m-frames of sequence m have a lower weight than all I-frame PDUs of sequences m to $m+L$. In the example of FIG. 3, $m =$

0 and $L = 3$. The length of the weighting interval, i.e. the number of sequences in the interval, is given by $L+1$, i.e. 4 in the example of FIG. 3.

The weighting interval can also be described for the PDUs of the less important P_l frames. Since in the example of FIG. 3 the P_l-frame PDUs of the

5   first sequence ($SQN = 0$) have assigned lower weights than the I-frame PDUs of the fifth sequence ($SQN = 4$) but higher weights that the I-frame PDUs of the sixth sequence ($SQN = 5$), the length of the weighting interval for the less important P_l-frame PDUs is equal $L+2$, thus exceeding the weighting interval length of the more important P_m-frame PDUs by one

10 sequence.

By using the weighting intervals as described above, the number of times that an I-frame PDU can be retransmitted before expiring is increased so that the probability of receiving the I-frame is also increased. This leads to transmissions or retransmissions of P_m-frame PDUs being delayed or

15 discarded in favor of I-frame PDU transmissions or retransmissions for a period of time that corresponds to $L+1$ sequences. Transmissions or retransmissions of P_m-frame PDUs of the first sequence do therefore not have highest priority up to the I-frame PDUs of the fourth sequence and these have time only up to their expiration, i.e. start of the sixth sequence, to

20 be successfully transmitted. The same applies to the P_l-frame PDUs although they have less time to be successfully transmitted or retransmitted to the receiver, as they are less important.

According to an embodiment, the length of the weighting interval is set according to the expiration time and the time-to-send of the PDU. For

25 example, the parameter L can be calculated based on the expiration time t_{exp}, the time-to-send t_{send}, the sequence duration T_{seq}, and the number of priority groups G:

$$L = \frac{t_{exp} - t_{send}}{T_{seq}} - G + 1$$

The value of T_{seq} denotes the duration of a sequence and is assumed to be fixed for all sequences.

According to an embodiment, a unique priority level is assigned to each video data unit, and the priority level is used for scheduling the data for
5 transmission. The embodiment overcomes the limitations of the prior art by considering not only frames scheduled for retransmission but also frames scheduled for first, i.e. regular, transmission. Furthermore, the embodiment assigns unique priority levels to the data thereby avoiding that multiple data units have the same priority. This is advantageous because at any time it is
10 clearly defined what data unit is to be scheduled for transmission.

According to the embodiment, the video data units that are available for transmission are selected (i.e. scheduled) according to their priority levels. This might lead to a situation where a less important frame is delayed or even discarded in favor of a more important frame, no matter of its position in
15 the stream. This increases the probability to receive this more important frame as it increases the amount of time for possible retransmissions of this frame.

Further, the embodiment is advantageous in that it allows for efficiently handling the available bandwidth for both, regular transmissions and
20 retransmissions. Regular transmissions and retransmissions share the same available bandwidth of the link. To which part the bandwidth is used by first time transmissions or by retransmissions is not fixed, but depends on the priority levels that are assigned to the data units, and on how many and which data units are requested for retransmission. By using the priority level
25 assignment according to the embodiments, the channel is utilized in an optimum way.

Moreover, the embodiment is advantageous in that it takes into account that the quality of the received compressed video stream depends on the degree in which the video quality is kept constant, without sudden improvements or
30 gradations of the actual seen video stream.

The assignment of a unique priority level is preferably done based on meta information of the video frame, e.g. the time-to-send, the time-to-display, the transmission delay, the importance or the type of compression, that is whether the frame is an I-frame or a P-frame. Through the use of this information and additional information on the current channel conditions, e.g. available bandwidth and receiver statistics, the embodiment allows for better deciding which frames are to be sent, re-sent or discarded.

Further, the embodiment provides a weighting algorithm that considers not only one single PDU but a weighting interval of multiple PDUs of multiple sequences (this will be discussed in more detail below). The weighting mechanism is advantageous since large fluctuations in the link quality do not result in large fluctuations of the quality of the received video stream.

While the invention has been described with respect to the physical embodiments constructed in accordance therewith, it will be apparent to those skilled in the art that various modifications, variations and improvements of the present invention may be made in the light of the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention. In addition, those areas in which it is believed that those of ordinary skill in the art are familiar, have not been described herein in order to not unnecessarily obscure the invention described herein. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrative embodiments, but only by the scope of the appended claims.

What is claimed is:

CLAIMS

1. An apparatus for transmitting video data to a video client, said video data being a series of video data units, the apparatus comprising:

   a priority assignment unit for assigning a unique priority level to each video data unit to be transmitted;

   a transmission buffer for buffering a plurality of video data units each having assigned a priority level, a first part of said plurality of video data units including video data units that have not yet been transmitted, a second part of said plurality of video data units including video data units that are available for retransmission;

   a transmission judgement unit for selecting, from said first and second parts, one of the buffered video data units according to its priority level;

   transmission means for transmitting the selected video data unit to the video client.

2. The apparatus of claim 1 wherein said video data units are packet datagram units.

3. The apparatus of claim 1 wherein said video data is transmitted in a video stream including intra-coded and predictive-coded frames, and said video data units are intra-coded or predictive-coded frames or segments thereof.

4. The apparatus of claim 1 further comprising a feedback evaluation unit for receiving feedback messages from the video client and operating said transmission judgment unit for performing retransmissions of video data units, wherein said transmission judgment unit is arranged for selecting said one of the buffered data units independently on whether the data unit is to be transmitted the first time or is to be retransmitted.

5. The apparatus of claim 1 wherein said priority assignment unit is arranged for assigning said unique priority level in accordance with the time-to-send and/or time-to-display and/or transmission delay of the respective video data unit.

5   6. The apparatus of claim 1 wherein said transmission judgment unit is triggered to select the video data unit in response to an acknowledgment or non-acknowledgment message from the video client.

7. The apparatus of claim 1 wherein said transmission judgment unit is triggered to select the video data unit whenever said transmission means

10   indicates that one of the previously selected video data units has been transmitted.

8. The apparatus of claim 1 wherein said transmission judgment unit is triggered to select the video data unit whenever said transmission buffer was empty and receives its first video data unit.

15   9. The apparatus of claim 1 wherein said video data is transmitted in a video stream including intra-coded frames and predictive-coded frames, and the predictive-coded frames are classified into at least two different priority classes.

10. The apparatus of claim 9 wherein the priority levels assigned to the intra-

20   coded frames and the priority levels assigned to the predictive-coded frames of one of the priority classes differ by an amount that depends on the overall maximum number of packet datagram units in the priority class, and/or on the number of priority classes.

11. The apparatus of claim 9 wherein the priority levels assigned to the

25   predictive-coded frames of one priority class and the priority levels assigned to the predictive-coded frames of another priority class differ by an amount that depends on the overall maximum number of packet datagram units in the priority class, and/or on the number of priority classes.

12. The apparatus of claim 1 wherein said video data units are arranged in sequences, and said priority assignment unit is arranged for assigning the priority levels to the video data units such that the priority levels gradually decrease from sequence to sequence.

13. The apparatus of claim 12 wherein said video data is transmitted in a video stream including intra-coded frames and predictive-coded frames, and the difference of the intra-coded frame priority levels and the predictive-coded frame priority levels depend on a parameter that indicates the length of a weighting interval, wherein said weighting interval is a group of several sequences in which the predictive-coded frame video data units of the first sequence in the weighting interval have a lower priority level than the intra-coded frame video data units of the last sequence in the weighting interval.

14. A method of transmitting video data to a video client, said video data being a series of video data units, the method comprising:

assigning a unique priority level to each video data unit to be transmitted;

buffering one or more video data units each having assigned a priority level, wherein a first part of the video data units have not yet been transmitted, and a second part of the video data units are available for retransmission;

selecting one of the buffered video data units from the first and second parts according to its priority level; and

transmitting the selected video data unit to the video client.

15. The method of claim 14 wherein said video data units are packet datagram units.

16. The method of claim 14 wherein said video data is transmitted in a video stream including intra-coded and predictive-coded frames, and said video data units are intra-coded or predictive-coded frames or segments thereof.
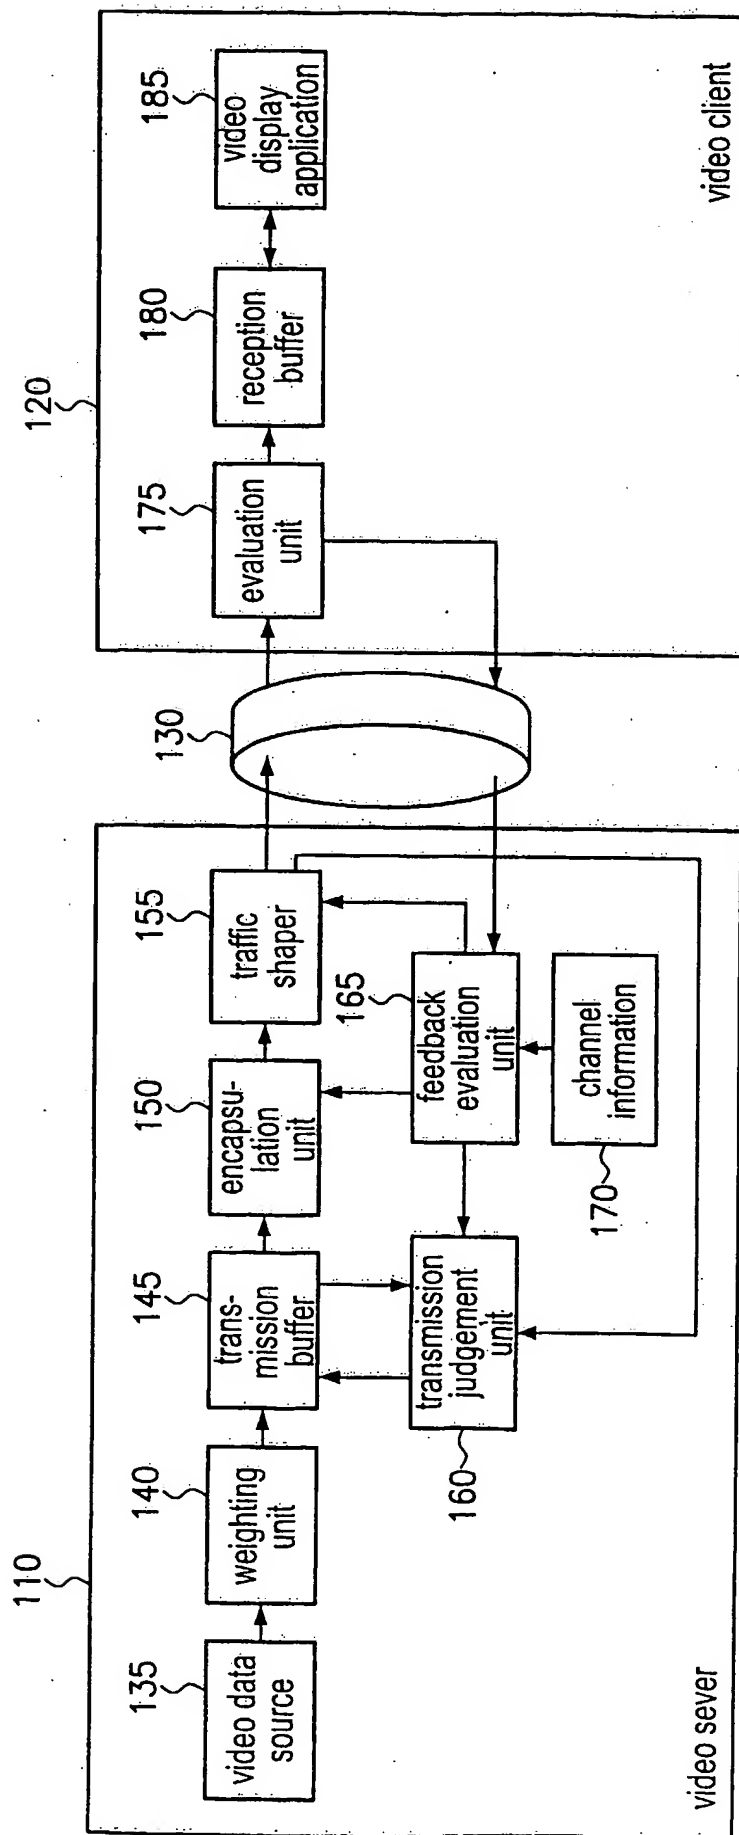
17.   The method of claim 14 further comprising a feedback evaluation
      mechanism comprising:

      receiving feedback messages from the video client; and

      performing retransmissions of video data units,

5   wherein said one of the buffered data units is selected independently on
      whether the data unit is to be transmitted the first time or is to be
      retransmitted.

18.   The method of claim 14 wherein assigning said unique priority level is
      performed in accordance with the time-to-send and/or time-to-display
10   and/or transmission delay of the respective video data unit.

19.   The method of claim 14 wherein said selection of the video data unit is
      triggered to be performed in response to an acknowledgment or non-
      acknowledgment message from the video client.

20.   The method of claim 14 wherein said selection of the video data unit is
15   triggered to be performed whenever one of the previously selected video
      data units has been transmitted.

21.   The method of claim 14 wherein said selection of the video data unit is
      triggered to be performed whenever a video data unit is buffered the first
      time.

20 22.   The method of claim 14 wherein said video data is transmitted in a video
      stream including intra-coded frames and predictive-coded frames, and the
      predictive-coded frames are classified into at least two different priority
      classes.

23.   The method of claim 22 wherein the priority levels assigned to the intra-
25   coded frames and the priority levels assigned to the predictive-coded
      frames of one of the priority classes differ by an amount that depends on

the overall maximum number of packet datagram units in the priority class, and/or on the number of priority classes.

24. The method of claim 22 wherein the priority levels assigned to the predictive-coded frames of one priority class and the priority levels assigned to the predictive-coded frames of another priority class differ by an amount that depends on the overall maximum number of packet datagram units in the priority class, and/or on the number of priority classes.

25. The method of claim 14 wherein said video data units are arranged in sequences, and said priority assignment is arranged for assigning the priority levels to the video data units such that the priority levels gradually decrease from sequence to sequence.

26. The method of claim 25 wherein said video data is transmitted in a video stream including intra-coded frames and predictive-coded frames, and the difference of the intra-coded frame priority levels and the predictive-coded frame priority levels depend on a parameter that indicates the length of a weighting interval, wherein said weighting interval is a group of several sequences in which the predictive-coded frame video data units of the first sequence in the weighting interval have a lower priority level than the intra-coded frame video data units of the last sequence in the weighting interval.
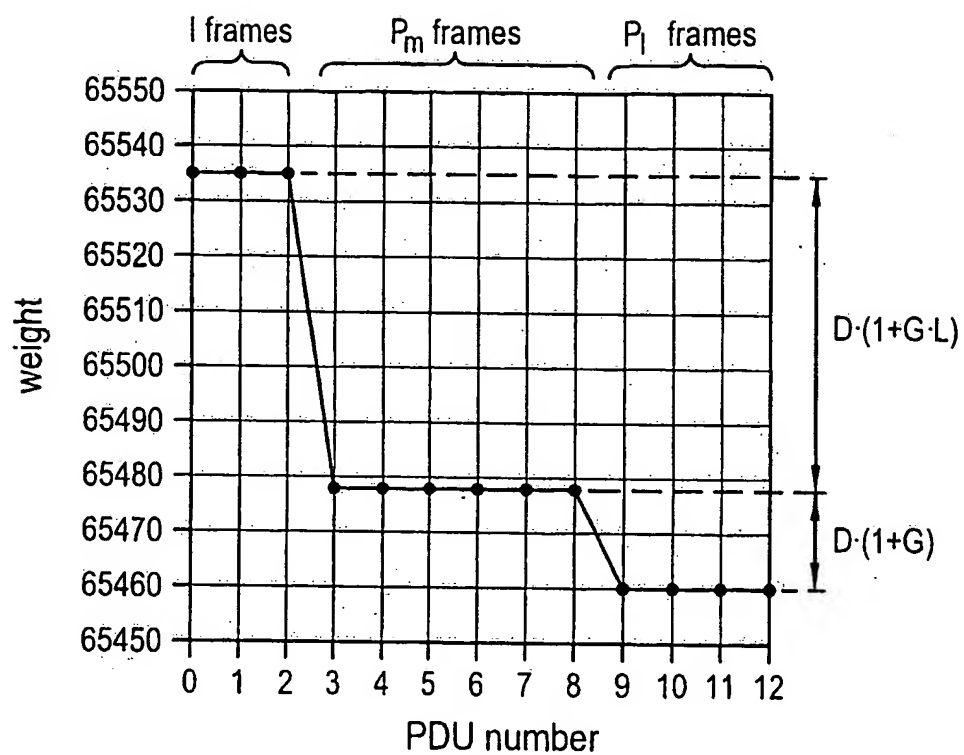
Fig. 1

FIG. 2

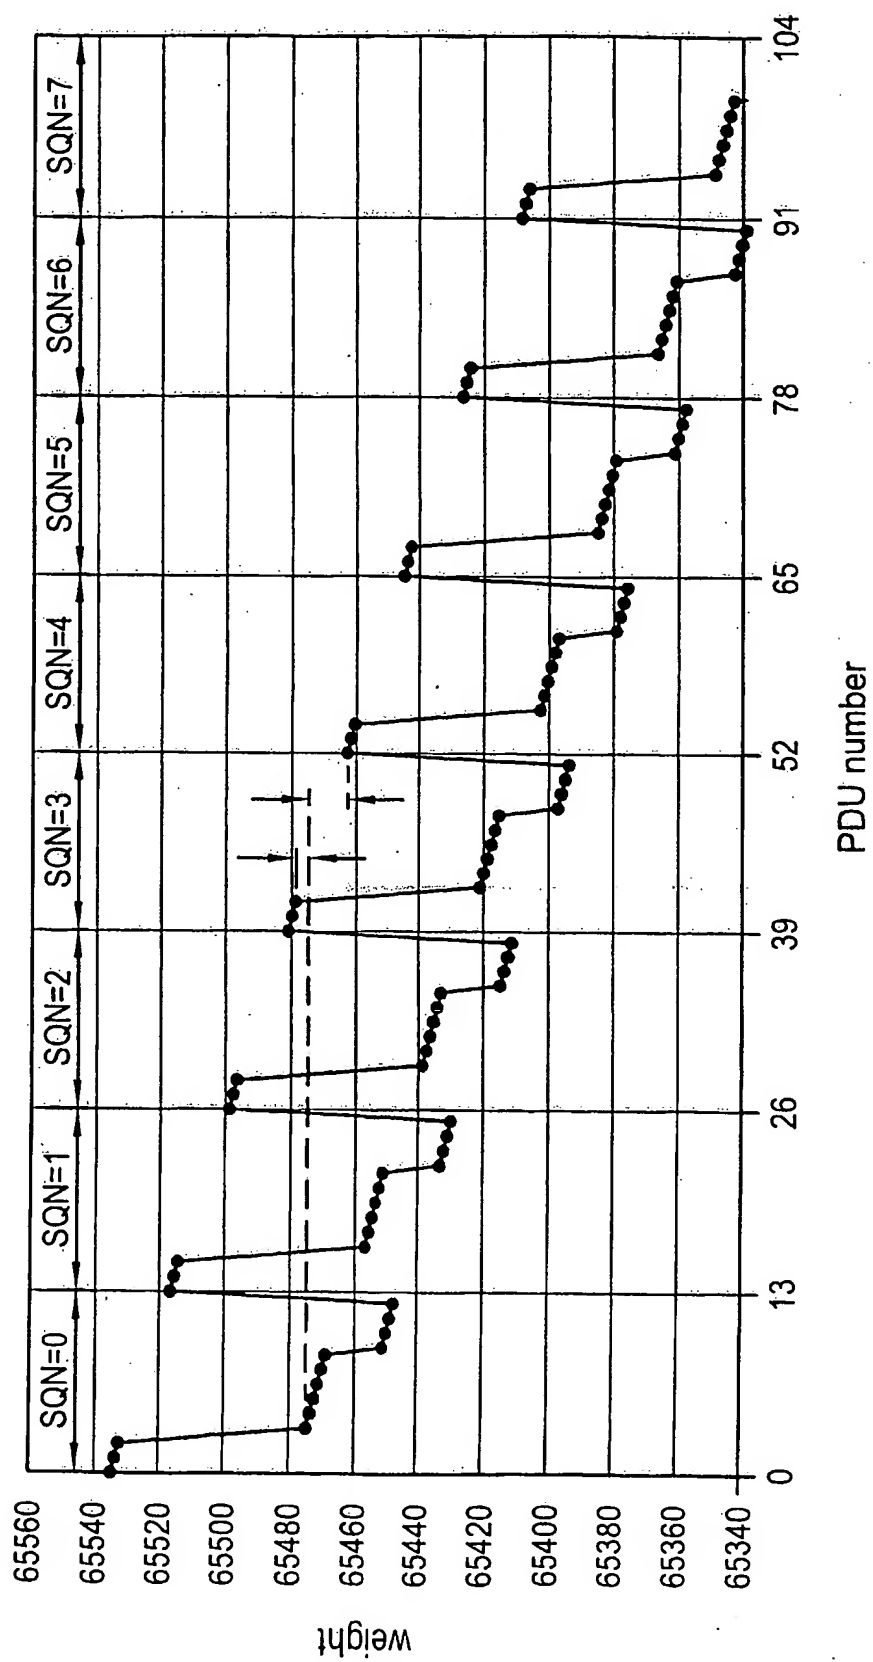FIG. 3